arXiv:1212.2747v1 [cs.LO] 12 Dec 2012

Bounds for the quantifier depth in two-variable logic

Christoph Berkholz
RWTH Aachen University
Institut für Informatik
D-52056 Aachen

Oleg Verbitsky*
Humboldt-Universität zu Berlin
Institut für Informatik
Unter den Linden 6
D-10099 Berlin

December 13, 2012

Abstract

We show examples of colored graphs G and H, both on n vertices, that in the existential two-variable logic can be distinguished only with quantifier depth more than $\frac{1}{11}n^2$. This bound is tight up to a constant factor. It contrasts with the fact that, if such graphs are distinguished in the unrestricted two-variable logic, they can be distinguished with quantifier depth at most n. We also discuss extensions of these results to logics with more variables and/or with more quantifier alternations. In particular, we show that the quantifier alternation hierarchy for k-variable logic is strict over colored trees. Without colors, this result holds true if $k \geq 3$, while the alternation hierarchy for two-variables logic collapses even if considered over all uncolored graphs.

1 Introduction

Let FO denote the class of all first-order formulas and FO^k be the subclass of formulas with occurrences of at most k variables. We will refer to FO^k as *k-variable logic*. The metatheory of these logics is rich in interesting results, even in the case of $k = 2$; e.g., [14, 15, 12, 13, 10, 31]. Finite-variable logics have found important applications in algorithmic model theory and descriptive complexity theory; see books [9, 21, 26] and survey articles [7, 14, 13, 19]. This is true even for "small" fragments of FO^k. The existential-positive k-variable logic is useful in heuristic analysis of the constraint satisfaction problem [25] and in estimating the complexity of resolution proofs [1, 3]. Propositional modal logic can be viewed as a fragment of FO^2; see [13]. The syntactic extension of FO^2 with counting quantifiers captures

---

*Supported by DFG grant VE 652/1–1. On leave from the Institute for Applied Problems of Mechanics and Mathematics, Lviv, Ukraine.

the known color-refinement procedure [22], an efficient way of graph canonization that works for almost all graphs.

We will consider a finite-variable logic over finite vertex-colored graphs (without directed edges), though our results for this class of structures carry over directly to directed vertex- and edge-colored graphs (i.e., to the structures over any vocabulary consisting of binary relations). Our aim is to investigate how expressibility of FO^k depends on the number of quantifier alternations allowed in a formula. In order to make a comparison, we use a quantitative measure of the expressive power of a logic $\mathcal{L} \subseteq \mathrm{FO}^k$, namely the quantifier depth needed to distinguish two given graphs in $\mathcal{L}$. This parameter is intimately related to the computational complexity of the logical equivalence problem for $\mathcal{L}$ [18, 2] and, hence, to testing isomorphism of graphs [6, 18]. Being a very natural descriptive complexity measure, this parameter appears also in other areas; for example, it corresponds to the depth of resolution refutations of width $k-1$ [1, 3].

More specifically, we say that a formula Φ *distinguishes* a graph G *from* a graph H if it is true on G but false on H. By $D_{\mathcal{L}}(G, H)$ we denote the minimum quantifier depth of such Φ over formulas in a logic $\mathcal{L}$; we write $D_{\mathcal{L}}(G, H) = \infty$ if G is indistinguishable from H in $\mathcal{L}$. We will simplify the notation by writing $D^k(G, H)$ in place of $D_{\mathrm{FO}^k}(G, H)$ and $D^k_{\mathcal{L}}(G, H)$ in place of $D_{\mathcal{L}\cap\mathrm{FO}^k}(G, H)$. Note that $D^k(G, H) = D^k(H, G)$.

Dawar, Lindell, and Weinstein [8] proved that $D^k(G, H) \le v(G)^k + k + 1$ if G and H are distinguishable in FO^k. Here and thoughout $v(G)$ denotes the number of vertices in the graph G. Immerman et al. [22, 6] discovered a connection between the parameter $D^k(G, H)$ and the number of color refinements needed for the k-dimentional Weisfeiler-Lehman procedure to distinguish G and H. By a careful analysis in [30], Oleg Pikhurko and the second author showed that this connection leads to a better bound, namely

$$D^k(G, H) \le v(G)^{k-1} \tag{1}$$

if $v(G) \ge 2$; see Section 3 where this bound is proved in detail for $k = 2$ (the treatment in [30] does not cover this case and states a little bit worse bound than (1)). If $k = 2$, this bound is tight up to a constant factor (for example, $D^2(P_n, P_{n+1}) \ge n/2$); whether this is so for $k \ge 3$ is a challenging question.

Bound (1) stays true for logic with *counting quantifiers*, for which we will use notation $D^k_{\#}(G, H)$. A linear lower bound $D^k_{\#}(G, H) = \Omega(v(G))$ is shown by Fürer [11] for each k (for an infinite sequence of pairs of graphs G and H with the same number of vertices).

The existential fragment $\Sigma_1 \cap \mathrm{FO}^k$ of k-variable logic consists of formulas that contain only existential quantifiers and are in the negation normal form (i.e., any negation stands in front of a relation symbol). For the existential logic we will often write $D^k_{\exists}(G, H)$ in place of $D^k_{\Sigma_1}(G, H)$. It is well known that $D^k_{\exists}(G, H) \le r$ iff Spoiler has a strategy allowing him to win the existential k-pebble Ehrenfeucht game on G and H in r rounds. If G is distinguishable from H in the existential k-variable logic, this equivalence immediately gives us the bound $D^k_{\exists}(G, H) \le (v(G)v(H))^k$, just because the number of possible positions in the game is bounded by $(v(G)v(H))^k$.

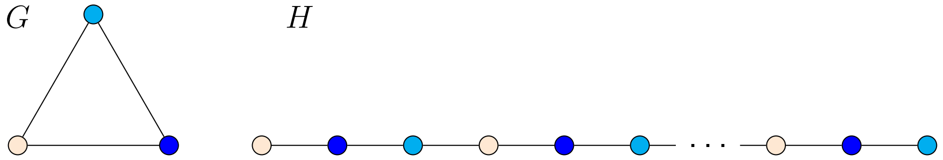

Figure 1: Spoiler has to make many passes through G because H is large.

Looking at the connection between logical distiguishability and the Ehrenfeucht game a bit more carefully, in Section 4 we prove a better bound

$$D^k_\exists(G,H) \le (v(G)v(H))^{k-1}. \tag{2}$$

This result is obtained by a rather general game-theoretic argument, that works as well for other logics having accompanying games, in particular, for the existential-positive k-variable logic.

The upper bound (2) for the existential case is worse than the best upper bound (1) for FO^k. One goal of our work is to show a gap between $D^k_\exists$ and D^k. As a first step it can be shown that $D^k_\exists(G,H)$ is lower bounded by $\Omega(v(H)^{k-1})$. Using the construction in [3] one can derive colored graphs G and H where Spoiler needs $\Omega(v(H)^{k-1})$ to win the existential k-pebble game. This observation is stated in the following theorem, whose proof is omitted here.

Theorem 1.1 *For every fixed $k \ge 2$ there is a colored graph G, a constant $c = c(k)$ and an infinite family $\mathcal{H}$ of colored graphs such that for all $H \in \mathcal{H}$*

$$c \cdot v(H)^{k-1} < D^k_\exists(G,H) < \infty. \tag{3}$$

For $k = 2$ a simple example is shown in Fig. 1. The case of $k \ge 3$ is, however, much more complicated. If compared to the upper bound (1), Theorem 1.1 shows a huge gap between the expressive power of k-variable logic and its existential fragment in terms of the parameter $D(G,H)$ if the graph G is small. However, Theorem 1.1 does not rule out the possibility that for graphs G and H with $v(G) = v(H) = n$ the difference disappears. We think that also in that case the upper bound (2) is tight, but given the current status of our research we have to state this as conjecture.

Conjecture 1.2 *For every fixed $k \ge 2$ there are infinite many pairs of colored graphs G and H with $v(G) = v(H) = n$ such that*

$$c \cdot n^{2k-2} < D^k_\exists(G,H) < \infty. \tag{4}$$

for a constant c depending on k only.

In this paper we settle the conjecture for the case $k = 2$. In Section 6 we construct infinitely many pairs of colored graphs G and H with $v(G) = v(H) = n$ such that G is distinguishable from H in the existential two-variable logic[1] with

$$D^2_\exists(G,H) > \frac{1}{11}\,n^2,$$

---

[1] and even in the existential-positive logic.

which matches the upper bound (2) up to a constant factor in the case of $k = 2$. One could wonder if it is necessary to consider colored graphs or if the lower bound also holds true for uncolored graphs. It turns out that the quantifier-rank of the existential two-variable logic drops down to linear over uncolored graphs. In Section 5 we prove that

$$D^2_{\exists}(G,H) \le 2\,v(H)$$

if G and H are uncolored. Even this particular case is non-obvious and requires a delicate analysis. On the other hand, our analysis shows that two-variable logic over uncolored graphs is much less complicated than over colored graphs: The distinguishability problem for this logic turns out to be decidable in logarithmic space in the former case (Corollary 4.3) and is known to be P-complete in the latter case (Kolaitis and Panttaja [23]).

Section 7 is devoted to higher levels of the quantifier alternation hierarchy. In the particular case of uncolored graphs we observe that, if such graphs G and H are distinguished in FO^2, then

$$D^2_{\Sigma_2}(G,H) \le v(G). \tag{5}$$

This shows that the alternation hierarchy over uncolored graphs collapses. The general case of colored graphs is also here much more complicated. Assuming that a colored graph G is distinguishable from a colored graph H in $\Sigma_2 \cap \mathrm{FO}^2$, we know neither if (5) stays true nor if there are G and H both with n vertices such that $D^2_{\Sigma_2}(G,H) = \Omega(n^2)$. At least, we are able to show that the alternation hierarchy of FO^2 over colored graphs is strict: For every $i \ge 1$ there are uncolored trees G_i and H_i such that G_i is indistinguishable from H_i in $\Pi_i \cap \mathrm{FO}^2$ while

$$D^2_{\Sigma_i}(G_i,H_i) \le i+1.$$

This results extends to the case of $k \ge 3$ variables, even over uncolored graphs.

Note that the quantifier alternation hierarchy has received a considerable attention in the past. Weis and Immerman [32] consider FO^2 over words in an alphabet A and show that the hierarchy collapses if A is finite and is strict if A is infinite. Bojańczyk and Segoufin [5] characterize tree languages definable in $\Sigma_2 \cap \Pi_2$ in terms of the descendant relation. Matz, Schweikardt and Thomas [27] establish that the alternation hierarchy of monadic second-order logic is strict over graphs, in fact, even over finite two-dimensional grids.

We conclude the paper with discussion of open questions in Section 8.

2 Preliminaries

A *(vertex-)colored graph* is a relational structure over vocabulary $\sigma = (E, U_1, U_2, \ldots)$, where E denotes a binary, irreflexive and symmetric, *adjacency relation* and U_i denotes a unary *color relation*. We will sometimes write $x \sim y$ instead of $E(x,y)$.

A first-order formula Φ over vocabulary σ *distinguishes a graph G from a graph H* if Φ is true on G but false on H. Note that, if Φ distinguishes G from H, then $\neg\Phi$ distinguishes H from G. Furthermore, we say that Φ distinguishes G *and*

H if it distinguishes G from H or H from G. A *logic* $\mathcal{L}$ is a class of formulas. For example, the *existential logic* consists of formulas in the negation normal form (where any negation stands in front of a relation symbol) with only existential quantifiers. By $D_{\mathcal{L}}(G,H)$ we denote the minimum quantifier depth of a formula $\Phi \in \mathcal{L}$ distinguishing G from H. In order to make this notation always meaningful, we write $D_{\mathcal{L}}(G,H) = \infty$ if G is indistinguishable from H by any formula in $\mathcal{L}$. For the existential logic we use notation $D_{\exists}(G,H)$. Moreover, the superscript k in $D^k_{\mathcal{L}}(G,H)$ will indicate that the k-variable fragment of $\mathcal{L}$ is considered. The absence of the subscript will mean that the full first-order logic over vocabulary σ is considered (i.e., $\mathcal{L} = \mathrm{FO}$).

The vertex set of a graph G is denoted by $V(G)$ and is supposed to be nonempty and finite. The number of vertices in G will be denoted by $v(G) = |V(G)|$. The *complement* of G is the graph on the same vertex set $V(G)$ with any two vertices adjacent if and only if they are not adjacent in G.

The *k-pebble Ehrenfeucht game on graphs G and H*, denoted by $\textsc{Ehr}^k(G,H)$, is played by two players, Spoiler and Duplicator, to whom we will refer as he and she respectively. The players have equal sets of k pairwise different pebbles. A *round* consists of a move of Spoiler followed by a move of Duplicator. Spoiler takes a pebble and puts it on a vertex in G or in H. Then Duplicator has to put her copy of this pebble on a vertex of the other graph. Duplicator's objective is to keep the following condition true after each round: the pebbling should determine a partial isomorphism between G and H. The variant of the k-pebble Ehrenfeucht game where Spoiler pebbles vertices always in G is called *existential* and denoted by $\textsc{Ehr}^k_{\exists}(G,H)$.

For each positive integer r, the r-round $\textsc{Ehr}^k(G,H)$ (as well as $\textsc{Ehr}^k_{\exists}(G,H)$) is a two-person game of perfect information with a finite number of positions. By Zermelo's theorem, either Spoiler or Duplicator has a *winning strategy* in this game, that is, a strategy winning against every strategy of the opponent.

Lemma 2.1 *Let $s \ge 1$. Then $D^2_{\exists}(G,H) \le s$ iff Spoiler has a winning strategy in the s-round game* $\textsc{Ehr}^2_{\exists}(G,H)$.

This well-known game characterization of existential logic generalizes to any finite-variable logic with bounded number of quantifier alternations; cf. [28].

3 Upper bound for the quantifier depth in FO^2

Our definition of a colored graph does not exclude that a vertex can have none or more than one colors. It will be convenient to regard the set of the colors that a vertex u has as a new color $C(u)$ (the mix of basic colors $U_1, U_2, \ldots$). The empty set of basic colors is allowed; thus, all originally uncolored vertices get colored in the same new color.

We will use a variant of the well-known color refinement procedure. Given a graph G, define a sequence of colorings of $V(G)$ as follows:

$$\begin{aligned} C^0(u) &= C(u), \\ C^{i+1}(u) &= \left\langle C^i(u), \left\{ C^i(a) : a \sim u \right\}, \left\{ C^i(a) : a \not\sim u \right\} \right\rangle . \end{aligned}$$

The following fact is a direct consequence of the connection between logical distinguishability and the color refinement procedure discovered in [6, 22].

Lemma 3.1 *Let $s \ge 1$. Then $D^2(G,H) \le s$ iff*

$$\left\{ C^{s-1}(u) : u \in V(G) \right\} \ne \left\{ C^{s-1}(v) : v \in V(H) \right\}. \tag{6}$$

Theorem 3.2 *Let G and H be colored graphs, both having at least two vertices. If they are distinguishable in two-variable logic, then*

$$D^2(G,H) \le \min \{v(G), v(H)\}.$$

Since $D^2(G,H) = D^2(H,G)$, it suffices to prove that $D^2(G,H) \le v(G)$. The proof is based on the well-known "color stabilization" argument; see [6, 22] or [30, Section 4]. Given a colored graph G, let C^i_G denote the partition of $V(G)$ into the monochromatic classes of the coloring C^i. Since $C^{i+1}(u) = C^{i+1}(u')$ implies $C^i(u) = C^i(u')$, the partition C^{i+1}_G is a refinement of C^i_G. Refinement cannot always be proper and eventually stabilizes: for some $i < v(G)$ we have $C^i_G = C^{i+1}_G$ and, hence, $C^i_G = C^j_G$ for all $j \ge i$. Call an index $s \ge 1$ *stable for G* if $C^{s-1}_G = C^s_G$.

Lemma 3.3 *If s is stable for G, then $D^2(G,H) \le s+1$ for any H distinguishable from G in the two-variable logic.*

Proof. We will need the following general notion. Given a partition P of $V(G)$, the *quotient graph* G/P has P as the vertex set with $A, B \in P$ being adjacent if G contains an edge between a vertex in A and a vertex in B. Note that G/P can have loops.

If Inequality (6) is true, then $D^2(G,H) \le s$ by Lemma 3.1. Suppose, therefore, that the two sets in (6) are equal.

If s is stable for H too, consider the stable partitions $P = C^{s-1}_G$ and $Q = C^{s-1}_H$. Consider also the one-to-one map $f : P \to Q$ taking each element of P to the element of Q of the same C^{s-1}-color. Note that f cannot be an isomorphism between the quotient graphs G/P and H/Q for else the color sets in (6) would be equal not only for the index s but also for all subsequent indices, and we would have $D^2(G,H) = \infty$ by Lemma 3.1. Therefore, P contains partition elements A and B that are adjacent in G/P while $f(A)$ and $f(B)$ are non-adjacent in H/Q or vice versa ($f(A)$ and $f(B)$ are adjacent while A and B are not). This implies that A and $f(A)$ have different C^s-colors. If $C \in P$ and $C \ne A$, then A and $f(C)$ also have different C^s-colors just because their C^{s-1}-colors are different. It follows that

$$\{ C^s(u) : u \in V(G)\} \ne \{ C^s(v) : v \in V(H)\} \tag{7}$$

and we have $D^2(G,H) \le s+1$ by Lemma 3.1.

It remains the case that s is not stable for H. Inequality (7) holds true because the set in the right hand side contains more colors. Thus, $D^2(G,H) \le s+1$ in both cases. ∎

The color stabilization argument readily implies that $n = v(G)$ is stable for any graph G. An improvement to $n-1$ takes some work.

Lemma 3.4 *The index $n-1$ is stable for every colored graph G on $n \geq 2$ vertices.*

Proof. If in the coloring C of G there are at least two colors, then $|C_G^0| \geq 2$, i.e., the partition C_G^0 consists of two elements, and we are done because then at most $n-2$ proper refinement steps are possible. Assume, therefore, that there is only one color or, equivalently, G is an uncolored graph.

If $n = 2$, then G is either complete or empty. By this reason $C_G^0 = C_G^1$, and the lemma is true. Assume, therefore, that $n \geq 3$.

We will need some graph-theoretic terminology. A set of vertices M is called a *module* of G if any two vertices in M have the same adjacency pattern to the rest of the graph, that is, if $x, y \in M$ and $z \notin M$, then $x \sim z$ iff $y \sim z$. Two vertices x and y are called *twins* if they form a module or, equivalently, if the transposition of x and y is an automorphism of the graph. A vertex is *isolated* if it is not adjacent to any vertex. A vertex is *universal* if it is adjacent to all other vertices. Furthermore, we will call a vertex *normal*, if it is neither isolated nor universal.

We will say that C_G^{i+1} is a *small refinement* of C_G^i if $|C_G^{i+1}| = |C_G^i| + 1$, that is, all but one elements of the partition C_G^i stay the same in C_G^{i+1} and the exceptional one is split up into two parts.

Claim A. If $C_G^0, C_G^1, \ldots, C_G^t$ is a sequence of t small refinements and $X \in C_G^{t-1}$ is split into Y and Z in C_G^t, then one of the two parts of X, say Y, consists of twins of G, and the other part, Z, is a module of G.

Proof of Claim A. We proceed by induction on t. In the base case, when $t = 1$, we have C_G^1 being a partition of $X = V(G)$ into two parts, Y and Z. Note that in the coloring C^1 there can appear three colors; let us call them I, U, and N. More specifically, all isolated vertices get the color I, all universal get U, and all normal get N. By assumption, exactly two of the three colors appear in C^1. Obviously, the combination $\{I, U\}$ is impossible. Suppose that the colors appearing in C^1 are I and N (the case of $\{U, N\}$ is similar; it is transformed to the case of $\{I, N\}$ by taking the complement of G). Let Y and Z consist of the isolated and normal vertices respectively. The claim is true because all isolated vertices are twins, and all non-isolated vertices form a module.

Let $t > 1$. By the induction assumption, all but one elements of C_G^{t-1} are twin-classes, and the exceptional element is a module. Since no twin-class can be split up, X is this exceptional element. Since X is a module, $\{Y, Z\}$ coincides with the partition $C_{G[X]}^1$, where $G[X]$ denotes the subgraph of G induced on X. By the argument used in the base case, Y consists either of isolated or of universal vertices of $G[X]$, and the remaining part Z is a module of $G[X]$. Referring once again to the fact that X is a module, we conclude that Y is a twin-class of G and Z is a module of G. $\triangleleft$

Consider now the sequence of partitions $C_G^0, C_G^1, \ldots, C_G^{n-2}$. If it contains a pair $C_G^i = C_G^{i+1}$, then $C_G^{n-3} = C_G^{n-2} = C_G^{n-1}$, and the lemma is true. If there is a pair

C_G^i, C_G^{i+1} with $|C_G^{i+1}| \geq |C_G^i| + 2$, then we are done by the color stabilization argument. In the only remaining case, the whole sequence consists of small refinements. Claim A applies to each $t \leq n-2$. If we have $|Y| \geq 2$ for some $t \leq n-2$, we are done because the twin-class Y cannot be split up later, and hence proper refinement terminates at the step $n-2$ at latest. Assume that $|Y| = 1$ for every $t \leq n-2$. Then for $t = n-2$ we have $|Z| = 2$, and Z is a module of G. Therefore, Z consists of two twins. It follows that $C_G^{n-1} = C_G^{n-2}$, completing the proof of the lemma. ■

Remark 3.5 Lemma 3.4 cannot be further improved because the value of $n-1$ is best possible for all graphs in the following sequence: G_1 consists of two isolated vertices, and G_{i+1} is obtained by taking the complement of G_i and adding an isolated vertex to it.

Theorem 3.2 follows directly from Lemmas 3.3 and 3.4.

We conclude this subsection with a brief discussion of the general case of k variables. The approach used here to prove Theorem 3.2 is borrowed from [30], where we focused on the case of $k \geq 3$. In [30] we observed that $D^k(G,H) < \mathit{Stab}^{k-1}(G)+k$, where $\mathit{Stab}^d(G)$ denotes the minimum stable index for the Weisfeiler-Lehman coloring of $V(G)^d$. If $d \geq 2$, this coloring is from the very beginning not monochromatic. More specifically, in the beginning $V(G)^d$ is partitioned into the types of d-tuples with respect to equality and adjacency relations. It follows that $\mathit{Stab}^d(G) \leq n^d - |C^0| + 1$, where $|C^0|$ denotes the size of the initial coloring of $V(G)^d$. If $n \geq d$, the number of equality types is equal to the Bell number B_d, which suffices to derive the bound (1) if $k \geq 4$. If $k = 3$, counting the equality types alone does not suffice but then it is enough to notice that, unless G is complete or empty, all but one equality types are split further by taking adjacency into account.

4 Upper bound for existential logic: Uncolored graphs

Theorem 4.1 *Let G and H be uncolored graphs. If G is distinguishable from H in the existential two-variable logic, then*

$$D^2_\exists(G,H) \leq 2\,v(H).$$

The proof of Theorem 4.1 takes the rest of this subsection. We will actively exploit the following graph-theoretic notions.

We call a graph *normal* if all its vertices are normal (i.e., it has neither isolated nor universal vertex). Note that a graph is normal iff its complement is normal. Let K be an induced subgraph of a graph G. If K is a normal graph, we call it a *normal subgraph* of G. Furthermore, we define a class of *abnormal graphs of rank i* inductively as follows:

- disconnected abnormal graphs of rank 1 are exactly empty graphs with at least 2 vertices;

- a disconnected abnormal graph of rank $i+1$ is a graph that consists of a connected abnormal graph of rank i and a number of isolated vertices;
- connected abnormal graphs of rank i are exactly complements of disconnected abnormal graphs of rank i.

The rank of an abnormal graph G will be denoted by $\operatorname{rk} G$. Note that examples of abnormal graphs already appeared in Remark 3.5 above.

Lemma 4.2 *Let G be a graph with $n \geq 2$ vertices. Then G has no normal subgraph iff G is abnormal.*

Proof. In order to prove the lemma in the forward direction, we use induction on n. Since G itself is not normal, it has either an isolated or an universal vertex. W.l.o.g., assume the former (otherwise consider $\overline{G}$). If all vertices in G are isolated, which is so in the base case of $n = 2$, this is an abnormal graph of rank 1. Otherwise, let G' be obtained from G by removing all isolated vertices. G' is an induced subgraph of G and, hence, has no normal subgraph. By the induction assumption, G' is abnormal. By construction, G' has no isolated vertex. Moreover, G' is connected for else it would have also no universal vertex and, hence, would be a normal subgraph of G. Therefore, G is abnormal with $\operatorname{rk} G = \operatorname{rk} G' + 1$.

For the backward direction, we use induction on $i = \operatorname{rk} G$. In the base case of $i = 1$, the graph G is either empty or complete and does not have any normal subgraph indeed. Let $i > 1$. Assume that G is disconnected (otherwise consider $\overline{G}$). Then G consists of an abnormal graph G' of rank $i-1$ and some isolated vertices. If an induced subgraph K of G contains at least one of the latter, this is an isolated vertex in K, and K is not a normal graph. Otherwise, K is an induced subgraph of G' and cannot be a normal graph by the induction assumption. ■

Suppose that G is an abnormal graph. We now introduce a ranking of vertices in G. If G is disconnected, it has at least one isolated vertex; if G is connected, there is at least one universal vertex. Denote the set of such vertices by ∂G. Every vertex in ∂G is assigned *rank* 1. In particular, if $\operatorname{rk} G = 1$, then all vertices of G have rank 1. If $u \notin \partial G$, then it is assigned *rank* one greater than the rank of u in the graph $G - \partial G$. The rank of a vertex u in G will be denoted by $\operatorname{rk} u$. It ranges from the lowest value 1 to the highest value $\operatorname{rk} G$.

Note that a vertex u has the same adjacency to all other vertices of equal or higher rank. If u is adjacent to all such vertices, we say that u is of *universal type*; otherwise we say that it is of *isolated type*. Note that the type gets flipped when the rank changes by one.

Let us turn to Theorem 4.1. By Lemma 2.1, we have to show that if Spoiler has a winning strategy in $\text{EHR}^2_{\exists}(G, H)$ for some number of rounds, then he is able to win fast enough, that is, in at most $2\,v(H)$ rounds. We split our analysis into several cases.

Case I: H has a normal subgraph K. Duplicator can win any number of rounds by playing inside K. Whatever vertex $v \in V(K)$ is pebbled in the preceding round, she is able to find another vertex $v' \in V(K)$, no matter whether it has to be adjacent to v or not.

From now on we will assume that H has no normal subgraph. By Lemma 4.2, H is abnormal.

Case II: G has a normal subgraph K while H is an abnormal graph. Spoiler has a strategy allowing him to win within $\mathrm{rk}\, H + 1 \le v(H)$ rounds. He always pebbles vertices inside K. Suppose that a pair of vertices $(u, v) \in V(K) \times V(H)$ is pebbled in the preceding round. If v is of universal type, Spoiler pebbles $a \in V(K)$ nonadjacent to u. Duplicator is forced to respond with a vertex b nonadjacent to v, which implies that $\mathrm{rk}\, b < \mathrm{rk}\, v$. If v is of isolated type, Spoiler pebbles $a \in V(K)$ adjacent to u and again forces pebbling $b \in V(H)$ of lower rank. If Duplicator respects adjacency all the time, Spoiler eventually forces $\mathrm{rk}\, b = 1$. He plays the next round in the same way, demanding the adjacency condition opposite to the adjacency that b has to the rest of the graph H, and wins. We will refer to this kind of strategy, that forces pebbling a vertex of always lower rank, as a *force-out strategy*.

From now on we will assume that, like H, the graph G has no normal subgraph. By Lemma 4.2, G is also abnormal.

Case III: Both G and H are abnormal. This case is most complicated and requires a finer analysis.

Claim B. Suppose that the players pebble a pair of vertices $(u, v) \in V(G) \times V(H)$. If $\mathrm{rk}\, u > \mathrm{rk}\, v$ or if $\mathrm{rk}\, u = \mathrm{rk}\, v$ and these vertices are of different types (i.e., one is of isolated type while the other is of universal type), then Spoiler is able to win no later than in the next $\mathrm{rk}\, v$ moves.

Proof of Claim B. Spoiler has a force-out strategy. We describe it recursively and simultaneously prove the claim by induction on $\mathrm{rk}\, v$.

If $\mathrm{rk}\, v = 1$, the vertex v is either isolated or universal in H. Note that the latter case is obtained from the former by considering the game on $\overline{G}$ and $\overline{H}$, which is virtually the same. Suppose, therefore, that v is isolated. If $\mathrm{rk}\, u > 1$, the vertex u is not isolated in G. If $\mathrm{rk}\, u = 1$, then u is not isolated as well (by assumption, it is universal in G). In the next round, Spoiler pebbles a vertex a adjacent to u and wins.

Let $\mathrm{rk}\, v \ge 2$. Assume first that u and v are of the same type. Thus, $\mathrm{rk}\, u > \mathrm{rk}\, v$. W.l.o.g., suppose that both u and v are of isolated type. In the next round Spoiler pebbles a vertex a of rank $\mathrm{rk}\, u - 1$. Since a is of universal type, it is adjacent to u. Since v is nonadjacent to all vertices of equal or heigher rank, Duplicator is forced to pebble a vertex b in H of rank $\mathrm{rk}\, b < \mathrm{rk}\, v$. Since $\mathrm{rk}\, b < \mathrm{rk}\, a$, Spoiler wins in the next $\mathrm{rk}\, b \le \mathrm{rk}\, v - 1$ moves by the induction assumption.

Assume now that u and v are of different types. W.l.o.g., suppose that u is of isolated type while v is of universal type. If $\mathrm{rk}\, u < \mathrm{rk}\, G$, Spoiler pebbles a vertex a of rank $\mathrm{rk}\, u + 1$. If $\mathrm{rk}\, u = \mathrm{rk}\, G$, Spoiler pebbles another vertex a of the same rank $\mathrm{rk}\, G$ (there are always at least two vertices of the highest rank). In either case, a is nonadjacent to u. This forces Duplicator to pebble a vertex b of rank $\mathrm{rk}\, b < \mathrm{rk}\, v$. Now $\mathrm{rk}\, b < \mathrm{rk}\, a$ and, again by the induction assumption, Spoiler wins in the next $\mathrm{rk}\, b \le \mathrm{rk}\, v - 1$ moves. $\triangleleft$

Case III-a: $\operatorname{rk} G > \operatorname{rk} H$ *or* $\operatorname{rk} G = \operatorname{rk} H$ *with one graph connected and the other disconnected.* Spoiler is able to win within $\operatorname{rk} H + 1 \leq v(H)$ rounds. In the first round, he pebbles a vertex u of rank $\operatorname{rk} H + 1$ or of the highest rank $\operatorname{rk} H$ if $\operatorname{rk} G = \operatorname{rk} H$. Whatever Duplicator's response v is, its rank is at most $\operatorname{rk} H$. Since the pair (u, v) meets the conditions of Claim B, Spoiler wins in the next $\operatorname{rk} H$ moves.

Case III-b: $\operatorname{rk} G = \operatorname{rk} H = r$ *and the graphs are both connected or both disconnected.*

We will say that G is *distinguishable from* H *at rank* ℓ if G has at least two vertices of rank ℓ while H has only one vertex of this rank. We split Case III-b further depending on whether or not such ℓ exists.

Case III-b-1: G *is indistinguishable from* H *at any rank* ℓ*.* Duplicator has a simple winning strategy, no matter how long the game is played. She always responds with a vertex having the same rank as the vertex pebbled in this round by Spoiler.

Case III-b-2: G *is distinguishable from* H *at some rank* ℓ*.* Spoiler has a strategy allowing him to win within $2r + 2 \leq 2\,v(H)$ rounds. In the first two rounds he pebbles two vertices of rank ℓ. Since there is only one vertex of this rank in H, this forces pebbling a pair $(u, v) \in V(G) \times V(H)$ such that $\operatorname{rk} u = \ell$ while $\operatorname{rk} v \neq \ell$. If $\operatorname{rk} v < \ell$, Spoiler applies the force-out strategy of Claim B and wins in at most $\operatorname{rk} v \leq \ell - 1 \leq r - 1$ extra moves. If $\operatorname{rk} v > \ell$, then v has higher rank than u. Spoiler's objective is now to increase the rank of a vertex pebbled in G, making it greater than the rank of the corresponding vertex in H; this will make the force-out strategy applicable again. This objective is achieved by the *overtaking strategy* described in the proof of the following claim.

Claim C. Once a pair (u, v) with $\operatorname{rk} u < \operatorname{rk} v$ is pebbled, Spoiler is able in at most $r - \operatorname{rk} v + 2$ moves to force pebbling a pair (u', v') such that $\operatorname{rk} u' > \operatorname{rk} v'$.

Proof of Claim C. Consider first the case that u and v are of different types. W.l.o.g., let u be of isolated type and v be of universal type. Spoiler pebbles a vertex u' of rank $\operatorname{rk} v$. Note that u and u' are not adjacent. Since in H all vertices of higher or equal rank are adjacent to v, Duplicator is forced to pebble a vertex v' of rank less than $\operatorname{rk} v = \operatorname{rk} u'$.

Now, consider the case that u and v are of the same type. Note that then $\operatorname{rk} u \leq \operatorname{rk} v - 2$. W.l.o.g., suppose that both u and v are of isolated type. In the next round, Spoiler pebbles a vertex a of rank $\operatorname{rk} v - 1$. Note that a and u are not adjacent, and a is of universal type. Let b denote Duplicator's response.

If $\operatorname{rk} b < \operatorname{rk} a$, Spoiler's objective is achieved. If $\operatorname{rk} b = \operatorname{rk} a = \operatorname{rk} v - 1$, then b is of universal type and, hence, adjacent to v; Duplicator is dead. Let $\operatorname{rk} b = \operatorname{rk} a + 1 = \operatorname{rk} v$. Though we still have $\operatorname{rk} a < \operatorname{rk} b$, the vertices a and b are now of different types, and Spoiler overtakes Duplicator in one extra move. Finally, assume that $\operatorname{rk} b > \operatorname{rk} v$. If a and b are of different types, Spoiler needs only one extra round as above. If a and b are of the same type, Spoiler plays the next round in the same way. It remains to notice that $\operatorname{rk} b$ can increase at most $r - \operatorname{rk} v$ times. $\triangleleft$

Let us turn to analysis of Subcase III-b-2. If $\operatorname{rk} u < \operatorname{rk} v$, then Spoiler applies the overtaking strategy. By Claim C, in at most $r - \operatorname{rk} v + 2 \leq r$ moves he forces pebbling a pair (u', v') with $\operatorname{rk} u' > \operatorname{rk} v'$ and then wins in at most $\operatorname{rk} v' \le r$ extra moves by using the force-out strategy.

Case III-c: $\operatorname{rk} G = \operatorname{rk} H - 1 = r$. W.l.o.g., suppose that H is disconnected. Let H' denote the graph obtained from H by removal of all isolated vertices. Thus, H' is a connected abnormal graph of rank r.

Case III-c-1: G is connected (while H is not). This case is essentially equivalent to Case III-b for the graphs G and H'. If G is indistinguishable from H' at any rank, then Duplicator survives in $\textsc{Ehr}^2_\exists(G, H')$ arbitrarily long, and the same holds true for $\textsc{Ehr}^2_\exists(G, H)$ (Duplicator just plays always in H').

Assume that G is distinguishable from H' at some rank ℓ. Like as in Case III-b-2, Spoiler wins $\textsc{Ehr}^2_\exists(G, H')$ in at most $2r + 2$ moves. Applying the same strategy in $\textsc{Ehr}^2_\exists(G, H)$, within the same number of rounds Spoiler either also wins (if Duplicator plays inside H') or forces pebbling a pair of vertices $(u, v) \in V(G) \times V(H)$ such that v is isolated in H (once Duplicator steps outside H'). In the latter case Spoiler pebbles a neighbor of u and wins. He can always do so, as G has no isolated vertex. Alltogether, Spoiler needs at most $2r + 3 < 2\,v(H)$ rounds to win.

Case III-c-2: G is disconnected (like H). Duplicator has a strategy allowing her resist arbitrarily long.

Suppose that Spoiler pebbles a vertex u. We say that Duplicator's move v in response is *straight* if $\operatorname{rk} v = \operatorname{rk} u$. We call it *shifted* if $\operatorname{rk} v = \operatorname{rk} u + 2$ (hence $\operatorname{rk} u < \operatorname{rk} G$) or if $\operatorname{rk} v = \operatorname{rk} u + 1$ while $\operatorname{rk} u = \operatorname{rk} G$. The pair of vertices (u, v) will be called *straight* or *shifted* correspondingly.

In the first round Duplicator makes a straight move. In general, suppose that a pair (u, v) was pebbled in the preceding round, and Spoiler pebbles $a \in V(G)$ in the current round. We will describe how Duplicator chooses her response $b \in V(H)$.

Assume that (u, v) is a straight pair. Whenever possible, Duplicator's move b is also straight. Note that making another straight move immediately after a straight move is never loosing for Duplicator. Playing so is impossible in the only case when $a \ne u$, $\operatorname{rk} u = \operatorname{rk} v$, and G is distinguishable from H at this rank. In this case Duplicator makes a shifted move, which is nonlosing also in this case because u and v are of the same type while both $\operatorname{rk} a \ge \operatorname{rk} u$ and $\operatorname{rk} b \ge \operatorname{rk} v$, hence $a \sim u$ iff $b \sim v$.

Assume now that (u, v) is a shifted pair. If $\operatorname{rk} a = \operatorname{rk} u$ and $a \ne u$, Duplicator makes a straight move. She does not lose because a and b are of the same type while both $\operatorname{rk} u \ge \operatorname{rk} a$ and $\operatorname{rk} v \ge \operatorname{rk} b$. If $\operatorname{rk} a \ne \operatorname{rk} u$, Duplicator makes a shifted move and does not lose. Indeed, w.l.o.g. suppose that $\operatorname{rk} a > \operatorname{rk} u$. Then $\operatorname{rk} v = \operatorname{rk} u + 2$ and $\operatorname{rk} b \ge \operatorname{rk} v$ (even when $\operatorname{rk} b = \operatorname{rk} a + 1$). Since u and v are of the same type, we have $a \sim u$ iff $b \sim v$.

Case III-d: $\operatorname{rk} G \le \operatorname{rk} H - 2$. Duplicator has a winning strategy for any number of rounds because H has an induced abnormal subgraph H' of $\operatorname{rk} H' = \operatorname{rk} G + 1$ and of the same connectivity as G. Duplicator always plays in H', following her winning strategy in $\textsc{Ehr}^2_\exists(G, H')$ as described above in Case III-c-2.

The proof of Theorem 4.1 is complete. As a by-product, it gives us also the following result.

Corollary 4.3 *The problem of deciding whether or not one uncolored graph is distiguishable from another in the existential two-variable logic is solvable in logarithmic space.*

Proof. Given uncolored graphs G and H, we actually have to decide if there is r such that Spoiler has a winning strategy in the r-round game $\textsc{Ehr}_{\exists}^{2}(G,H)$. It suffices to determine which of the cases I–III and which of their subcases takes place for G and H. This task easily reduces to the following problem: Given a graph G, we have to recognize if it is abnormal and, if so, we should be able to compute rank and type of each vertex of the graph. It turns out that this is not much harder than computing vertex degrees.

Note first that the degree of a vertex of isolated type increases together with its rank. The degree of a vertex of universal type decreases as its rank increases. Moreover, any vertex of universal type has larger degree than any vertex of isolated type. It suffices to check this inequality between vertices of different types that have the two highest possible ranks. Let $\operatorname{rk} u = \operatorname{rk} G$ and $\operatorname{rk} v = \operatorname{rk} G - 1$ and suppose that u is of isolated type. Then the degree of u is equal to the number U of vertices of universal type in G while the degree of v is one smaller that the sum of U and the number R of all vertices of the highest rank $\operatorname{rk} G$. If u is of universal type, then $\deg u = U - 1$ while $\deg v = U - R$. Since $R \geq 2$, we have $\deg v > \deg u$ in the former case and $\deg u > \deg v$ in the latter case.

These observations suggest the following algorithm for recognition of abnormal graphs and, in passing, for computing ranks of vertices in an abnormal graphs and determining their types. If an input graph G has neither isolated nor universal vertex, it is not abnormal. If G has a universal vertex, it suffices to check if $\overline{G}$ is abnormal. Suppose, therefore, that G has an isolated vertex. We first compute all vertex degrees. Suppose that we obtain d different degree values. Classify a vertex as being of isolated type if its degree is among $\frac{d+1}{2}$ smallest values and as being of universal type otherwise. Assign odd ranks to vertices of isolated type and even ranks to vertices of universal type by the following rules: The larger degree a vertex has, the higher rank it receives if it is of isolated type and the lower rank it receives if it is of universal type. Finally, we have to check if the following adjacency condition is obeyed: Vertices of isolated type are nonadjacent to all vertices of equal or higher rank, while vertices of universal type are adjacent to all vertices of equal or higher rank (in particular, vertices of the same degree are twins). ∎

Remark 4.4 The analysis of Case II in the proof of Theorem 4.1 leads to a monochromatic analog of the example shown in Fig. 1. Let H_n be an abnormal graph of rank $n-1$ on n vertices. Fix an arbitrary normal graph, for example, P_4. Then $D_{\exists}^{2}(P_4, H_n) = n$.

5 Upper bound for existential logic: Colored graphs

Theorem 5.1 *Let G and H be colored graphs. If G is distinguishable from H in the existential two-variable logic, then*

$$D^2_\exists(G,H) \le v(G)v(H).$$

Proof. By Lemma 2.1, we have to prove that, if Spoiler has a winning strategy in the s-round game $\textsc{Ehr}^2_\exists(G,H)$ for some s, then he has a winning strategy in the game with $v(G)v(H)$ rounds.

The proof is based on a general game-theoretic argument. Consider a two-person game, where the players follow some fixed strategies and one of them wins. Then the length of the game cannot exceed the total number of all possible positions because once a position occurs twice, the play falls into an endless loop. Here it is assumed that players's strategies are *non-adaptive*, that is, that a *strategy* of a player maps a current *position* (rather than the whole sequence of previous positions) to one of the positions available for the player. We have to show that we really can use this assumption for the Ehrenfeucht game. Moreover, since the number of positions in $\textsc{Ehr}^2_\exists(G,H)$ is $(v(G)v(H))^2$, we have to "reduce" it. To this end, instead of the space $v(G)^2 \times v(H)^2$ of positions in the 2-pebble Ehrenfeucht game, we will consider the space $v(G) \times v(H)$ of *configurations*. The idea appears aleady in the standard proof of Lemma 2.1: the players can play the 2-pebble Ehrenfeucht game optimally by using strategies that are non-adaptive in a rather strong sense. The strong non-adaptiveness means here that the next move is computed based on the outcome of the preceding round (which is only a part of the current position formed by the outcome of the two preceding rounds).

Given $u \in V(G)$ and $v \in V(H)$, let $r_{u,v}$ denote the smallest $r \ge 0$ such that Spoiler has a winning strategy in the r-round game $\textsc{Ehr}^2_\exists(G,u,H,v)$, where the vertex pair (u,v) is pebbled before the game starts. Here $r_{u,v} = 0$ if even the initial configuration (u,v) is already losing for Duplicator, that is, u and v have different colors. Furthermore, we fix a strategy $S_{u,v}$ winning for Spoiler in the $r_{u,v}$-round $\textsc{Ehr}^2_\exists(G,u,H,v)$.

Claim D. Suppose that $r_{u,v} \ge 1$. Let a be the vertex of G that $S_{u,v}$ brings forth in the first round of $\textsc{Ehr}^2_\exists(G,u,H,v)$. Then, for every $b \in V(H)$, we have $r_{a,b} < r_{u,v}$.

Proof of Claim D. If we had $r_{a,b} \ge r_{u,v}$, Duplicator could win the $r_{u,v}$-round $\textsc{Ehr}^2_\exists(G,u,H,v)$ by pebbling b in the first round. $\triangleleft$

Define a strategy S for Spoiler as follows.

1-ST ROUND. Spoiler pebbles a vertex $u_1 \in V(G)$ such that $\max_{w \in V(H)} r_{u,w}$ is as small as possible (we fix such u_1 arbitrarily if there is more than one choice).

$(i+1)$-TH ROUND. If the pair $(u,v) \in V(G) \times V(H)$ was pebbled in the i-th round, Spoiler places the other pebble on the vertex a produced by $S_{u,v}$ in the first round of $\textsc{Ehr}^2_\exists(G,u,H,v)$.

We now argue that S is a winning strategy for Spoiler in the $v(G)v(H)$-round game $\textsc{Ehr}^2_\exists(G,H)$. Indeed, let Spoiler follow S, fix a strategy for Duplicator arbitrarily, and denote the outcome of the i-th round by (u_i, v_i). Assume that Duplicator survives in t rounds and consider the sequence

$$(u_1, v_1), (u_2, v_2), \ldots, (u_t, v_t). \tag{8}$$

Thus, we have $r_{u_i,u_i} > 0$ for all $i \le t$ and, by Claim D, we also have $r_{u_i,u_i} > r_{u_{i+1},u_{i+1}}$ for all $i < t$. It follows that all pairs in the sequence (8) are pairwise different. If at least two colors appear in the graphs, then at least one pair (u, v) must have $r_{u,v} = 0$. Therefore, not all pairs from $V(G) \times V(H)$ can appear in the sequence (8), and we have $t \le v(G)v(H) - 1$, which implies the lemma.

Suppose now that all vertices G and H are equally colored. Since this is equivalent to the uncolored case, Theorem 4.1 applies and gives us $D^2(G,H) \le 2\,v(H)$, which stays whithin the bound $v(G)v(H)$ unless $v(G) = 1$. In the last case the theorem is trivially true just because the single-vertex graph G is indistinguishable from any H in $\mathrm{FO}^2_\exists$. ∎

Theorem 5.1 can easily be extended to any k-variable logic.

Theorem 5.2 *Let $k \ge 3$. If a graph G is distinguishable from a graph H in the existential k-variable logic, then*

$$D^k_\exists(G,H) < (v(G)v(H))^{k-1}.$$

Proof. In order to adapt the proof of Theorem 5.1 for the general case of k variables, we consider tuples of vertices $u \in V(G)^{k-1}$ and $v \in V(H)^{k-1}$ and again define $r_{u,v}$ to be the smallest r such that Spoiler has a winning strategy in the r-round Ehrenfeucht game $\textsc{Ehr}^2_\exists(G,u,H,v)$ on G and H with the initial configuration $(u,v) \times V(G)^{k-1} \times V(H)^{k-1}$. We can suppose that in the first round Spoiler puts all k pebbles on the same vertex. Assume that the $(i-1)$-th round has been played with k-tuples of vertices $\bar{u} \in V(G)^k$ and $\bar{v} \in V(H)^k$ being pebbled. Given $j \le k$, let $\sigma_j \bar{u} \in V(G)^{k-1}$ denote the $(k-1)$-tuple obtained from $\bar{u}$ by removal of the j-th coordinate. In the i-th round, the processs of making decision by Spoiler can be split into two steps.

1. He first decides which of the k pairs of pebbles will play in the next round.
2. If the j-th pair is chosen, then Spoiler actually plays the first round of the game $\textsc{Ehr}^2_\exists(G,u_i,H,v_i)$, where $u_i = \sigma_j \bar{u}$ and $v_i = \sigma_j \bar{v}$.

Assume that Spoiler plays optimally, that is, in Step 1 he chooses j minimizing the value r_{u_i,v_i}, and in Step 2 he invokes a winning strategy for the r_{u_i,v_i}-round game $\textsc{Ehr}^2_\exists(G,u_i,H,v_i)$. Then the elements of the sequence (8) are pairwise different. This implies that $t \le (v(G)v(H))^{k-1} - 2$ because not all pairs from $V(G)^{k-1} \times V(H)^{k-1}$ can occur in (8); for example, none of the pairs (u,v) where u and v have different equality types. ∎

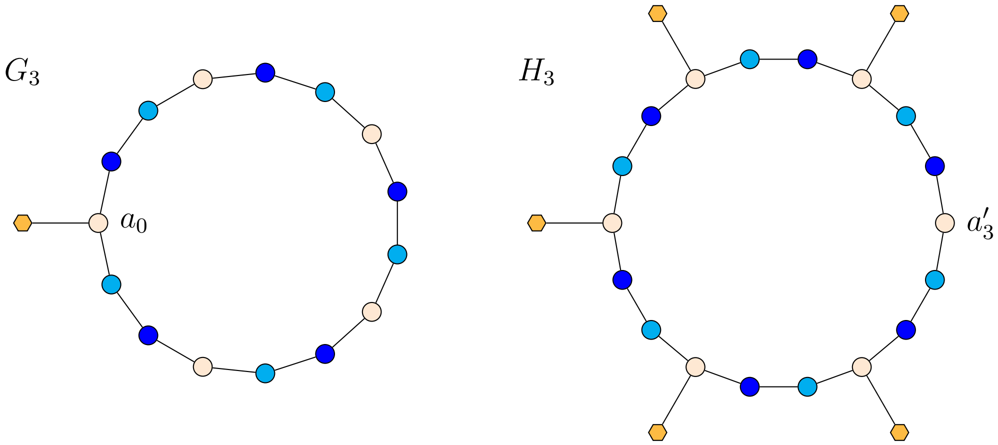

Figure 2: Proof of Theorem 6.1. The only way for Spoiler to win is to force pebbling the pair of vertices (a_0, a'_3) in a round. Duplicator can force Spoiler to make many passes around G_m before this goal is achieved.

6 Lower bound for the existential two-variable logic

Theorem 6.1 *There are infinitely many colored graphs G and H, both on n vertices, such that*

$$D^2_\exists(G,H) > \tfrac{1}{11}\, n^2$$

while G is distinguishble from H in the existential-positive two-variable logic.

Proof. Our construction will depend on an integer parameter $m \ge 2$. We construct a pair of colored graphs G_m and H_m such that G_m is distinguishable from H_m in the existential two-variable logic, both $v(G_m) = O(m)$ and $v(H_m) = O(m)$, and $D^2_\exists(G_m, H_m) = \Omega(m^2)$. Though $v(G_m) < v(H_m)$, later we will be able to increase the number of vertices in G_m to $v(H_m)$.

The graphs have vertices of 4 colors, namely apricot, blue, cyan, and dandelion. G_m contains a cycle of length $3(2m-1)$ where apricot, blue, and cyan alternate in this order; see Fig. 2. H_m contains a similar cycle of length $3 \cdot 2\,m$. Successive apricot, blue, and cyan vertices will be denoted by a_i, b_i, and c_i in G_m, where $0 \le i < 2m-1$, and by a'_i, b'_i, and c'_i in H_m, where $0 \le i \le 2m-1$. Furthermore, the vertex a_0 is adjacent to a dandelion vertex d_0, and every a'_i except for $i = m$ is adjacent to a dandelion vertex d'_i. This completes the description of the graphs.

By Lemma 2.1, we have to show that Spoiler is able to win $\textsc{Ehr}^2_\exists(G_m, H_m)$ and that Duplicator is able to prevent losing the game for $\Omega(m^2)$ rounds.

Note that, once the pair (a_0, a'_m) is pebbled, Spoiler wins in the next move by pebbling d_0. He is able to force pebbling (a_0, a'_m) as follows. In the first round he pebbles a_0. Suppose that Duplicator responds with a'_s, where $0 \le s < m$. In a series of subsequent moves, Spoiler goes around the whole circle in G_m, visiting $c_{2m-2}, b_{2m-2}, a_{2m-2}, c_{2m-2}, \ldots$ (if $m < s < 2m$, he does the same but in the other direction). As Spoiler comes back to a_0, Duplicator is forced to arrive at a'_{s+1}. The next Spoiler's tour around the circle brings Duplicator to a'_{s+2} and so forth. Thus, the most successful moves for Duplicator in the first round is a'_0. Then Spoiler needs to play $1 + m(2m - 1) + 1 = 2m^2 - m + 2$ rounds in order to win.

Since Spoiler keeps the adjacency relation true in every round, he wins even the *existential two-pebble game*, where Duplicator's objective is relaxed to keeping a partial *homomorphism*. It follows by [24] that G is distinguishble from H in the existential-positive two-variable logic.

Our next task is to design a strategy for Duplicator allowing her to survive $\Omega(m^2)$ rounds, no matter how Spoiler plays. The idea is actually similar to the idea beyond the simple example depicted by Fig. 1: Duplicator is able to force Spoiler to pass around the cycle in G_m many times. A crucial observation is that (a_0, a'_m) is the only pair whose pebbling allows Spoiler to win in one extra move.

Let us regard the additive group $\mathbb{Z}_{2m}$ as a cycle graph with i and j adjacent iff $i - j = \pm 1$. Denote the distance between vertices in this graph by Δ. The same letter will denote the following partial function $\Delta : V(G_m) \times V(H_m) \to \mathbb{Z}$. For two vertices of the same color, say, for a_i and a'_j, we set $\Delta(a_i, a'_j) = \Delta(i, j)$. Note that $\Delta(a_0, a'_m) = m$, which is the largest possible value. Duplicator' strategy will be to keep the value of the Δ-function on the pebbled pair as small as possible.

Specifically, in the first round Duplicator responds to Spoiler's move x with pebbling a vertex x' such that $\Delta(x, x') = 0$ (that is, if $x = a_i, b_i, c_i, d_0$, then $x' = a'_i, b'_i, c'_i, d'_0$ respectively). Suppose that a pair (y, y') is pebbled in the preceding round and Spoiler is still alive. If Spoiler pebbles x in the current round, Duplicator chooses her response x' by the following criteria.

- x' should have the same color as x and, moreover, $x' \sim y'$ iff $x \sim y$ (this is always possible unless $(y, y') = (a_0, a'_m)$ and $x = d_0$);
- if there is still more than one choice, x' should minimize the parameter $\Delta(x, x')$.

We do not consider the cases when $x = y$ or when x is pebbled by the pebble removed from y because, in our analysis, we can assume that Spoiler uses an *optimal* strategy, allowing him to win $\text{EHR}^2_{\exists}(G_m, y, H_m, y')$ in the smallest possible number of rounds (if he does not play optimally, Duplicator survives even longer).

Claim E. If $x \not\sim y$ and $x \neq y$, then $\Delta(x, x') \le 1$.

Proof of Claim E. Assume first that $x \neq d_0$ and $y \neq d_0$. W.l.o.g., suppose that y and y' are apricot and, specifically, $y' = a'_j$ (the blue and the cyan cases are symmetric to the apricot case). Not to lose immediately, Duplicator cannot pebble x' in $\{c'_{j-1}, a'_j, b'_j\}$, where $j - 1$ is supposed to be an element of $\mathbb{Z}_{2m}$. This can obstruct attaining $\Delta(x, x') = 0$ (if $x \in \{c_{j-1}, a_j, b_j\}$), but then there is a choice of x' with $\Delta(x, x') = 1$.

Assume now that $x = d_0$. Then $x' = d'_0$ if $y' \neq a'_0$ and $x' = d'_1$ otherwise. In both cases $\Delta(x, x') \le 1$.

Finally, let $y = d_0$ and $y' = d'_j$. Then the value $x' = a'_j$ is forbidden and, if this prevents $\Delta(x, x') = 0$, then we have $\Delta(x, x') = 1$. $\triangleleft$

Consider now the dynamical behaviour of $\Delta(x, x')$, assuming that Duplicator uses the above strategy and Spoiler follows an optimal winning strategy. We have

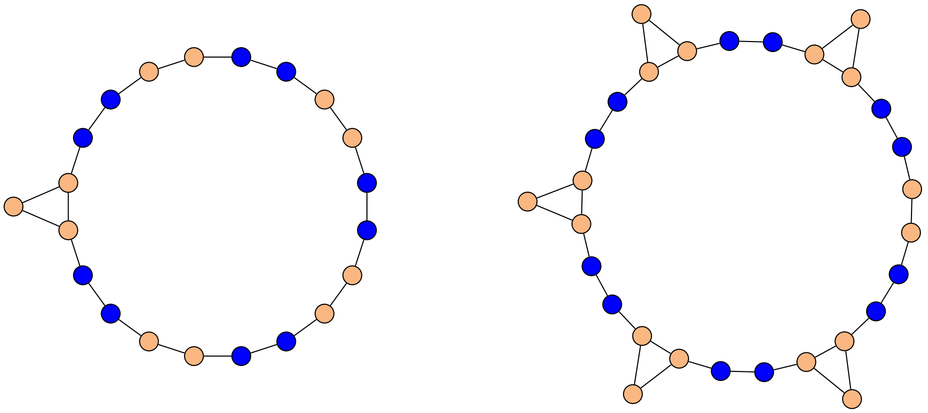

Figure 3: A bicolored modification of the example in Fig. 2.

$\Delta(x, x') = 0$ at the beginning of the game and $\Delta(x, x') = m$ at the end (that is, in the round immediately before Spoiler wins). Consider the last round of the game where $\Delta(x, x') \le 1$. By Claim E, starting from the next round Spoiler always moves along an edge in G_m. Note that, from now on, visiting d_0 earlier than in the very last round would be inoptimal. Therefore, Spoiler walks along the circle. Another consequence of optimality is that he moves always in the same direction.

By symmetry, we can suppose that Spoiler moves in the ascending order of indices. Note that $\Delta(x, x')$ increases by 1 only under the transition from $x = a_{2m-2}$ to $x = a_0$ (at point, the index of x makes a jump in $\mathbb{Z}_{2m}$, while the index of x' moves along $\mathbb{Z}_{2m}$ always continuously). In order to increase $\Delta(x, x')$ from 1 to m, the edge $a_{2m-2}a_0$ must be passed $m - 1$ times. It follows that, before Spoiler wins, the game lasts at least $2 + (m - 2) \cdot 3(2m - 1) = 6m^2 - 15m + 8$ rounds.

Note that $v(G_m) = 6m-2$ and $v(H_m) = 8m-1$. In order to make the number of vertices in both graphs $n = 8m-1$, let m be multiple of 3 and add two new connected components to G_m, namely the cycle of length $2m$ with alternating colors apricot, blue, and cyan and one isolated vertex af any color. Spoiler can still win by playing in the old component. Since playing in the new components does not help him, the game on the modified G_m and the same H_m lasts at least $6m^2-15m+8 = \frac{3}{32}n^2-O(n)$ rounds. ∎

Remark 6.2 In order to facilitate the exposition, the construction of graphs G_m and H_m uses 4 colors. In fact, the same idea can be realized with 2 colors. Appropriately modified G_m and H_m are shown in Fig. 3. The larger graph H_m has now $n = 10m - 1$ vertices, and the game lasts more than $\frac{2}{25}n^2 - O(n)$ rounds. For simplicity of presentation, we did not try to optimize neither the constant factor of $\frac{2}{25}$ here nor the factor of $\frac{1}{11}$ in Theorem 6.1. However, two colors are optimal in view of Theorem 4.1.

Remark 6.3 Theorem 6.1 has no analog for the logic with counting. Indeed, let G and H be distinguishable in $\mathrm{FO}^k_\#$ and $v(G) = v(H) = n$. Then $D^k_\#(G, H) = D^k_{\#,\exists}(G, H)$ because the universal quantifier can be replaced with the counting quantifier $\exists^n$. It remains to note that $D^2_\#(G, H) \le n$ similarly to Theorem 3.2.

7 Quantifier alternation hierarchy

Rephrasing Remark 6.3, one can say that the quantifier alternation hierarchy of $\mathrm{FO}^k_{\#}$ collapses over pairs of graphs with the same number of vertices. Let us look at the standard logic without counting.

For each $i \geq 1$, let Σ_i (resp. Π_i) denote the set of first-order formulas Φ in the negation normal form with the following property: Any sequence of nested quantifiers in Φ has at most $i-1$ quantifier alternations and begins with $\exists$ (resp. $\forall$) if it has exactly $i-1$ quantifier alternations. By definition, $\Sigma_i \cup \Pi_i \subset \Sigma_{i+1} \cap \Pi_{i+1}$ and Σ_1 is existential first-order logic. By the *quantifier alternation hierarchy* we mean the interlacing chains $\Sigma_1 \subset \Sigma_2 \subset$ and $\Pi_1 \subset \Pi_2 \subset$. We are interested in the corresponding fragments of a finite-variable logic. Sometimes we can write $D^k_{\exists\forall}(G,H)$ for $D^k_{\Sigma_2}(G,H)$ and $D^k_{\forall\exists}(G,H)$ for $D^k_{\Pi_2}(G,H)$.

If $k \geq 3$, we will see that the alternation hierarchy of FO^k is strict even over uncolored trees. If $k=2$, the colored and the uncolored cases are different. While the alternation hierarchy of FO^2 is still strict over colored trees, over uncolored graphs it collapses in the sense made precise in the following result.

Theorem 7.1 *If uncolored graphs G and H are distinguished in* FO^2*, then they are distinguished in its Σ_2-fragment with quantifier depth at most* $\max\{v(G), v(H)\}$.

Proof. We make use of our analysis done in the proof of Theorem 4.1. It gives us a complete description of the pairs of graphs with $D^2_{\exists}(G,H) < \infty$ and, by swapping G and H, also a complete description of the pairs with $D^2_{\forall}(G,H) < \infty$. In all these cases, exept III-b-2 and the similar Case III-c-1, we have

$$\min\{D^2_{\exists}(G,H), D^2_{\forall}(G,H)\} \leq \max\{v(G), v(H)\}.$$

In the two exceptional cases we have a twice worse bound because Spoiler may need to use the overtaking strategy before he runs the force-out strategy. However, the overtaking phase is not needed at all if Spoiler is allowed to jump to the other graph. This observation shows that we have

$$\min\{D^2_{\exists\forall}(G,H), D^2_{\forall\exists}(G,H)\} \leq \max\{v(G), v(H)\} \tag{9}$$

whenever $\min\{D^2_{\exists}(G,H), D^2_{\forall}(G,H)\} < \infty$.

The last condition is not fulfilled in two cases. First, this happens if both pairs (G,H) and (H,G) fall into Case III-b-1. In this case we actually have $D^2(G,H) = \infty$. The other situation when this happens is the subcase of Case I where both G and H have normal subgraphs. In the rest of the proof we show that also here we have the bound (9) unless $D^2(G,H) = \infty$.

It suffices to prove that, unless $D^2(G,H) = \infty$, the graphs G and H are distinguishable in $\Sigma_2 \cap \mathrm{FO}^2$ with quantifier depth at most $v(G)$. We proceed by induction on $v(G)$. Note that all graphs with less than 4 vertices are abnormal and, hence, $v(G) = 4$ in the base case.

If both G and H are normal, then $D^2(G,H) = \infty$. Note that this includes the base case because the only way for a graph on 4 vertices to have normal subgraphs is to be normal itself. If exactly one of G and H is normal, that is, the graphs are distinguished by the existence of an isolated or a universal vertex in one of them, then they are clearly distinguished in $\Sigma_2 \cap \mathrm{FO}^2$ with quantifier depth 2. Suppose, therefore, that none of G and H is normal. If one of them is connected and the other is not, then the graphs are again distinguished in $\Sigma_2 \cap \mathrm{FO}^2$ with quantifier depth 2. It remains to consider the case that G and H are both connected or both disconnected. W.l.o.g., assume the latter, i.e., that both G and H have isolated vertices.

If one of the graphs has a single isolated vertex and the other of them has at least two, then they are distinguished in $\Sigma_2 \cap \mathrm{FO}^2$ with quantifier depth 3 (Spoiler pebbles 2 isolated vertices in one of the graph and wins in the next round). It remains to consider the case when both G and H have exactly one isolated vertex or when both of them have at least two isolated vertices.

Remove all isolated vertices from G and H and denote the resulting graphs by G' and H'. Note that both G' and H' must contain normal subgraphs. By the induction assumption, either Duplicator survives in $\textsc{Ehr}^2(G', H')$ arbitrarily long or Spoiler needs at most $v(G')$ rounds to win the 1-jump version of $\textsc{Ehr}^2(G', H')$ where he is allowed to start the game in any graph at his discretion. In the former case, Duplicator survives arbitrarily long as well in $\textsc{Ehr}^2(G, H)$: whenever Spoiler pebbles an isolated vertex, Duplicator does the same in the other graph. In the latter case, Spoiler succeeds in the same type of the Ehrenfeucht game on G and H, needing perhaps one extra move to win. He plays inside G' and H', using his winning strategy in the game on these graphs. Playing always inside G' and H', Duplicator loses. Suppose that she pebbles an isolated vertex in G or H. Since Spoiler's counterpart of this vertex in H' or G' is not isolated, Spoiler wins in the next move by pebbling its neighbor. Note that no extra jump is required in this situation. By the induction assumption, Spoiler wins the 1-jump $\textsc{Ehr}^2(G, H)$ in $v(G') + 1 \le v(G)$ rounds. ■

Theorem 7.2 *For every $i \ge 1$ there are bicolored trees G_i and H_i such that*

1. $v(G_i) = v(H_i)$,
2. $D^2_{\Sigma_i}(G_i, H_i) \le i + 2$,
3. $D^2_{\Pi_i}(G_i, H_i) = \infty$.

The proof of Theorem 7.2 is a reminiscence of an earlier work on logical complexity of trees in [4, 29]. The same idea was used by Oleg Pikhurko (personal communication, 2007) to construct, for each i, a sequence of pairs of trees G_j and H_j such that $D_{\Sigma_i}(G_j, H_j) = O(1)$ while $D_{\Pi_i}(G_j, H_j) \to \infty$ as $j \to \infty$.

Proof. We use induction on i. The graphs G_1 and H_1, both being trees with 7 vertices colored in apricot or blue, are shown in Fig. 4. We have $D^2_\exists(G_1, H_1) \le 2$ because Spoiler wins $\textsc{Ehr}^2_\exists(G_1, H_1)$ by pebbling two blue vertices in G_1. We have $D^2_\forall(G_1, H_1) = \infty$ because Duplicator has a simple winning strategy in $\textsc{Ehr}^2_\exists(H_1, G_1)$.

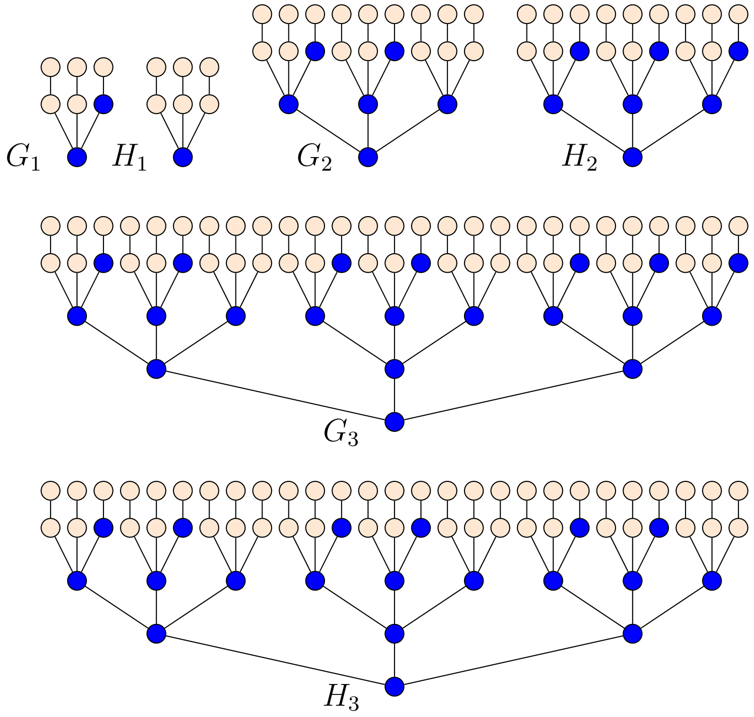

Figure 4: Proof of Theorem 7.2. The recursive construction of the sequence of graphs G_i and H_i.

We designate the vertices of degree 3 in G_1 and H_1 as *roots*. Both G_1 and H_1 have three *leaves* of degree 1. The other three vertices will be called *stalks*. In general, we will use the latter term for any vertex adjacent to a leaf. If we cut off the leaves, the truncated G_1 and H_1 still satisfy Conditions 1–3 of the theorem. However, this exra level of vertices is needed in the recursive construction of G_i and H_i for $i \geq 2$.

Suppose that the graphs G_{i-1} and H_{i-1}, rooted trees with the same number of vertices colored in apricot or blue, are already constructed. The graph H_i consists of three disjoint copies of G_{i-1} and a new vertex colored in blue, designated as root in H_i, which is adjacent to the root of each G_{i-1}-part. The graph G_i is constructed similarly but, instead of three G_{i-1}-branches, it has two G_{i-1}-branches and one H_{i-1}-branch; see Fig. 4. Obviously, G_i and H_i have equal number of vertices.

In order to prove the bound $D^2_{\Sigma_i}(G_i, H_i) \leq i + 2$, it suffices to design a strategy for Spoiler allowing him to win the Σ_i-version of $\textsc{Ehr}^2(G_i, H_i)$ in $i+2$ moves, where he begins playing in G_i and up to $i - 1$ times is allowed to jump from one graph to another. The following *climbing strategy* serves this purpose. Spoiler pebbles the root in the first round and then successively pebbles vertices along a path upwards arriving at a blue leaf (he is allowed to jump $i - 1$ times and continue such a path in the other graph). Note that exactly this strategy is successful for him in $\textsc{Ehr}^2_{\exists}(G_1, H_1)$, where Duplicator is forced either to break adjacency or to pebble an apricot stalk. Using induction in i, we show that Spoiler can force the same for every $i \geq 2$ in $i + 1$ moves under the condition that also Duplicator all the time climbs up in the trees.

In the first round, when Spoiler pebbles the root of G_i, Duplicator is forced to respond with pebbling the root of H_i. Otherwise, if she responds with a blue non-root vertex u in H_i, Spoiler jumps to H_i and pebbles the vertices along a path starting at u and finishing at an apricot stalk. Suppose that Duplicator does respond with pebbling the root of H_i.

In the second round Spoiler pebbles the root of the H_{i-1}-branch in G_i. Not to lose immediately, Duplicator responds with pebbling the root of one of the G_{i-1}-branches in H_i. Spoiler jumps to H_i and, starting from the third round, plays the Σ_{i-1}-game $\textsc{Ehr}^2(G_{i-1}, H_{i-1})$ using the climbing strategy. If Duplicator always climbs up, Spoiler wins in the next $i - 1$ rounds by the induction assumption.

There remains a possibility that Duplicator at some moment makes a move downwards. Then Spoiler, staying at a blue vertex w, starts pebbling the vertices along a shortest path from w to an apricot vertex (no extra jump is needed in this case). Duplicator loses because the distance from her standpoint to any apricot vertex is larger. Note that, if w is not a stalk, Spoiler wins by arriving at an apricot stalk; if w is a blue stalk, Spoiler pebbles the adjacent leaf. This is exactly the reason why we need an extra level of apricot vertices in G_1 and H_1.

It remains to prove that $D^2_{\Pi_i}(G_i, H_i) = \infty$. We describe a strategy for Duplicator in the Π_i-game $\textsc{Ehr}^2(G_i, H_i)$ (where Spoiler begins playing in H_i and is allowed to jump from one graph to another at most $i - 1$ times) and prove by induction on i that the strategy allows her to resist arbitrarily long. An important feature of this strategy is that Duplicator will always respects tree levels, that is, the distance to the root. Playing $\textsc{Ehr}^2_\forall(G_1, H_1)$, Duplicator is able to respect tree levels and vertex colors simultaneously, so we are done in the base case.

Let $i \geq 2$ and assume that Duplicator has a successful strategy for the Π_{i-1}-game $\textsc{Ehr}^2(G_{i-1}, H_{i-1})$. In the course of the game, we will call a branch of G_i or H_i *free* if it currently contains no pebbled vertex. We will call a branch of G_i and a branch of H_i *linked* if they contain vertices pebbled in a round. Duplicator plays the Π_i-game $\textsc{Ehr}^2(G_i, H_i)$ as follows.

- If Spoiler pebbles the root in G_i or H_i, Duplicator pebbles the root in the other graph.

- If Spoiler moves in a free G_{i-1}-branch of G_i or H_i, Duplicator mirrors his move in a free G_{i-1}-branch in the other graph. Note that this is always possible because both G_i and H_i have at least two G_{i-1}-branches.

- If Spoiler moves in the H_{i-1}-branch of G_i when it is free, Duplicator responds in a free G_{i-1}-branch of H_i according to her strategy in the Π_{i-1}-game $\textsc{Ehr}^2(G_{i-1}, H_{i-1})$.

- If Spoiler places a pebble in a branch containing a vertex pebbled by the other pebble, Duplicator responds in the linked branch. If the linked branches are both copies of G_{i-1}, Duplicator mirrors Spoiler's move; otherwise she follows her strategy in the Π_{i-1}-game $\textsc{Ehr}^2(G_{i-1}, H_{i-1})$.

Note that this strategy is level-preserving because, by assumption, so is the strategy on G_{i-1} and H_{i-1}.

Let us argue that Duplicator never loses the Π_i-game on G_i and H_i. The strategy description ensures that her move in the first round is non-losing. Suppose that pairs of vertices (u, v) and (a, b) are pebbled in two successive rounds. Assume first that u and v are the roots of G_i and H_i. To prove that Duplicator does not lose, we have to check that adjacency and colors are preserved. Indeed, we have $a \sim u$ iff $b \sim v$ because the strategy is level-preserving. Furthermore, a and b have the same color because, if they both are in G_{i-1}-branches, then Duplicator played by mirroring; while if a is in the H_{i-1}-branch of G_i, then Duplicator actually plays the first round of the game on G_{i-1} and H_{i-1} (using a non-losing strategy).

Assume now that u and v are non-roots. The case when Spoiler makes a move, a or b, in a free branch of G_i or H_i is tackled by the same argument as above. This case is even a little simpler because we clearly have $a \not\sim u$ and $b \not\sim v$. It remains to consider the case when u and a are in the same branch of G_i, and v and b are in the same branch of H_i. If these branches are copies of G_{i-1}, then adjacency and colors are preserved because Duplicator played by mirroring. If u and a are in the H_{i-1}-branch of G_i, then Duplicator is alive because she uses a non-losing strategy for $\mathrm{EHR}^2(G_{i-1}, H_{i-1})$. A crucial observation here is that the game on G_{i-1}- and H_{i-1}-subgraphs can start only after Spoiler jumps from G_i to H_i. As one jump is aleady used, this is a Π_{i-1} game, and the induction assumption does apply. ∎

Theorem 7.2 generalizes to any k-variable logic and, if $k > 2$, then no vertex coloring is needed any more.

Theorem 7.3 *Let $k \geq 3$. For every $i \geq 1$ there are uncolored trees G_i and H_i such that*

1. $v(G_i) = v(H_i)$,
2. $D^3_{\Sigma_i}(G_i, H_i) \leq i + 5$,
3. $D^k_{\Pi_i}(G_i, H_i) = \infty$.

Proof. If colors (apricot and blue) are allowed, then the construction of G_i and H_i from the proof of Theorem 7.2 can easily be adapted to suit k variables. Moreover, if there are more than two variables, the levels of apricot vertices in G_1 and H_1 are not needed any more. The only their purpose was to ensure Spoiler's win in the situation when the players arrived at vertices $u_1 \in V(G_i)$ and $v_1 \in V(H_i)$ immediately before the stalk level, after which Spoiler pebbled a blue stalk, say, $v_2 \in V(H_i)$, while Duplicator pebbled a blue vertex $u_2 \in V(G_2)$ in the lower level (because all stalks adjacent to u_1 were apricot). Now, that Spoiler has a third pebble, he puts it on v_3, the other blue neighbor of v_1. Since u_1 has only one blue neighbor, Duplicator loses.

For each $k \geq 3$, we now define trees G_i and H_i as follows. Let H_1 consist of a blue root-vertex and $k+1$ apricot leaves. G_1 is obtained from H_1 by recoloring one leaf in blue. Furthermore, H_i consists of the blue root and $k+1$ G_{i-1}-branches. G_i

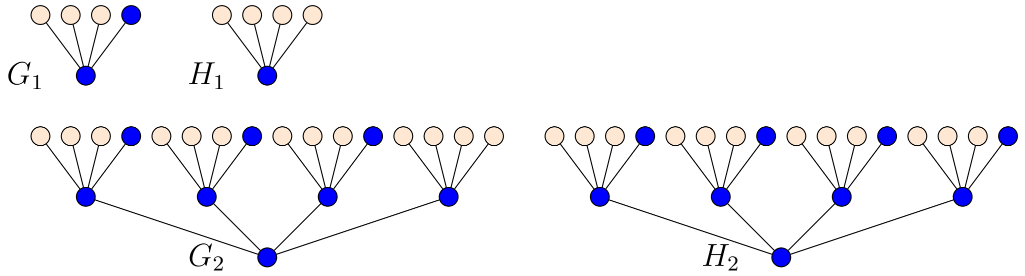

Figure 5: Proof of Theorem 7.3. The trees for 3-variable logic are still colored.

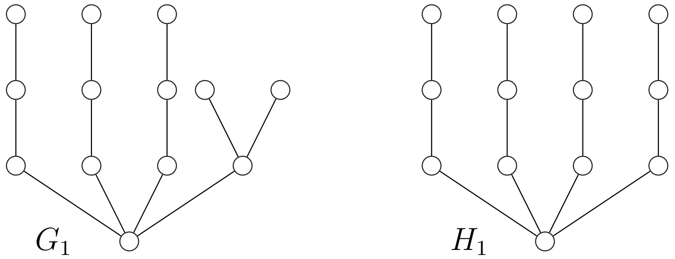

Figure 6: Proof of Theorem 7.3. The uncolored versions of G_1 and H_1 for 3-variable logic.

is obtained from H_i by replacing one G_{i-1}-branch with an H_{i-1}-branch; see Fig. 5. Except the aforementioned situation, analysis of Σ_i- and Π_i-games on G_i and H_i is completely similar to the proof of Theorem 7.2.

In order to remove colors from G_i and H_i, we grow two extra levels of vertices in these trees. Specifically, we connect each apricot leaf a to a new vertex a_1, which is adjacent to another new vertex a_2. Each blue leaf b is connected to two new vertices b_1 and b_2, which are non-adjacent to one another; see Fig. 6. Note that the modified graphs have vertices of four kinds: of degree 1 (all those are *new* vertices), of degree 2 (which are of two sorts: *old apricot* vertices and adjacent new vertices), and of degree at least 3 (*old blue* vertices).

In order to win the Σ_i-game on the modified uncolored graphs, Spoiler uses the climbing strategy as in the proof of Theorem 7.2. In at most $i + 2$ moves he forces pebbling a pair of old vertices that had different colors. Here we have two cases: Either Spoiler uses up all jumps and pebbles an old blue vertex while Duplicator pebbles an old apricot vertex (which includes the new situation discussed above), or Spoiler has at least one jump left over and pebbles an old apricot vertex while Duplicator pebbles an old blue vertex. Assume for a while that Spoiler can use 4 pebbles. In the former case he wins by pebbling 3 neighbors of the old blue vertex; in the latter case he makes a remaining jump and wins in the same way.

It remains to explain how Spoiler can win with 3 pebbles after an old blue vertex u (now of degree al least 3) and an old apricot vertex v (now of degree 2) are pebbled in a round. He first pebbles two neighbors u_1 and u_2 of u. Duplicator must respond with the two neighbors v_1 and v_2 of v. In the next round Spoiler moves the pebble from u to its third neighbor u_3. Duplicator must remove the pebble from v and place it on some vertex v_3 non-adjacent to both v_1 and v_2. Note that, while the distance between any two vertices of u_1, u_2, and u_3 equals 2, there is a pair of indices s and t such that u_s and u_t are at the distance more than 2. Spoiler now wins by moving the pebble from u_q to u, where $\{q\} = \{1, 2, 3\} \setminus \{s, t\}$.

Finally, we have to check that Duplicator survives arbitrarily long in the Π_i-game on the modified uncolored graphs. This can be shown by induction on i, much the same as in the proof of Theorem 7.2. ■

Corollary 7.4 *Let $k \geq 2$. The quantifier alternation hierarchy for FO^k over colored trees is strict: For every i there is a a class of colored trees $\mathcal{T}$ such that some sentence in $\Sigma_i \cap \mathrm{FO}^k$, if considered over colored trees, defines $\mathcal{T}$ while no sentence in $\Pi_i \cap \mathrm{FO}^k$ does it. Moreover, if $k \geq 3$, this holds true even for uncolored trees.*

Proof. Let G_i and H_i be colored trees as in Theorem 7.2 if $k = 2$ or uncolored trees as in Theorem 7.3 if $k \geq 3$. Let $\Phi \in \Sigma_i \cap \mathrm{FO}^k$ be a sentence distinguishing G_i from H_i. Define $\mathcal{T}$ to be the class of all trees satisfying Φ. Note that $G_i \in \mathcal{T}$ while $H_i \notin \mathcal{T}$. By this reason, no sentence in $\Pi_i \cap \mathrm{FO}^k$, if considered over trees, defines $\mathcal{T}$ (since it cannot distinguish G_i from H_i). ■

8 Questions and comments

We now state a few remaining problems, focusing on the "smallest" particular cases that are open.

Let G and H be colored graphs with $v(G) = v(H) = n$. We have $D^2(G, H) \leq D^2_{\exists\forall}(G, H) \leq D^2_{\exists}(G, H)$, where $D^2(G, H) \leq n$ if it is finite and $D^2_{\exists}(G, H)$ can be finite but large, namely $\Omega(n^2)$. Where is $D^2_{\exists\forall}(G, H)$ in this range (between $O(n)$ and $\Omega(n^2)$)?

Furthermore, is it true that there exists k, say $k = 3$, such that if $D^2(G, H) < \infty$, then we have $D^k(G, H) = \mathcal{O}(\log n)$? In other words, do we have an exponential speed-up in the quantifier depth if we increase the number of variables?

If we initially have 3 variables, a negative answer follows from a result of Immerman [20] who constructed pairs of colored graphs G and H such that $D^3(G, H) < \infty$ while $D_{\mathrm{FO}}(G, H) > 2^{\sqrt{\log n} - 1}$. However, the following modification of the question seems open: Do we have a speed-up under the stronger condition that G is *definable* in 3-variable logic, that is, $D^3(G, H) < \infty$ for *every* H non-isomorphic to G.

Questions of this kind are related to the graph isomorphism problem. Assume that there is a function f, for example $f(k) = 3k$, such that if a graph G on n vertices is definable in $\mathrm{FO}^k_{\#}$, then it is definable in $\mathrm{FO}^{f(k)}_{\#}$ with quantifier depth $O(\log n)$. Such a speed-up result would imply that isomorphism of graphs excluding a fixed minor could be tested in NC. This implication follows from two known facts. First, graphs excluding a fixed minor M are definable in $\mathrm{FO}^k_{\#}$ for some fixed $k = k(M)$ (Grohe [16, 17]), and hence distinguishability in this logic coincides with isomorphism within this class of graphs. Second, distinguishability with logarithmic quantifier depth is decidable in NC (Grohe and Verbitsky [18]).

The best lower bound for $D^k_{\#}(G,H)$, obtained by Fürer [11], is linear. In fact, Fürer's construction demonstrates a rich dynamical behavior. For each $h > 1$, he constructs a sequence of pairs of graphs G_n and H_n with $v(G_n) = \Theta(n)$ and $v(H_n) = \Theta(n)$ such that

$$D^k_{\#}(G_n, H_n) = \begin{cases} \infty & \text{if } k \le h, \\ \Theta(n) & \text{if } h < k \le 2h, \\ \Theta(\sqrt{n}) & \text{if } k = 2h+1, \\ O(\log n) & \text{if } k = 3h. \end{cases}$$

It is an interesting question whether such a behavior is "typical".

Acknowledgement

The authors are thankful to Oleg Pikhurko for his useful suggestions.

References

[1] A. Atserias and V. Dalmau. A combinatorial characterization of resolution width. *J. Comput. Syst. Sci.*, 74(3):323–334, 2008.

[2] C. Berkholz. Lower bounds for existential pebble games and k-consistency tests. In *Proceedings of the 27th Annual IEEE Symposium on Logic in Computer Science, LICS 2012, Dubrovnik, Croatia, June 25-28, 2012*, pages 25–34. IEEE, 2012.

[3] C. Berkholz. On the complexity of finding narrow proofs. In *Proc. FOCS'12*, 2012.

[4] T. Bohman, A. M. Frieze, T. Luczak, O. Pikhurko, C. D. Smyth, J. Spencer, and O. Verbitsky. First-order definability of trees and sparse random graphs. *Combinatorics, Probability & Computing*, 16(3):375–400, 2007.

[5] M. Bojanczyk and L. Segoufin. Tree languages defined in first-order logic with one quantifier alternation. *Logical Methods in Computer Science*, 6(4), 2010.

[6] J.-Y. Cai, M. Fürer, and N. Immerman. An optimal lower bound on the number of variables for graph identifications. *Combinatorica*, 12(4):389–410, 1992.

[7] A. Dawar. Finite models and finitely many variables. In D. e. a. Niwiński, editor, *Logic, algebra, and computer science. Helena Rasiowa in memoriam.*, volume 46 of *Banach Cent. Publ.*, pages 93–117. Polish Academy of Sciences, Institute of Mathematics, Warszawa, 1999.

[8] A. Dawar, S. Lindell, and S. Weinstein. Infinitary logic and inductive definability over finite structures. *Inf. Comput.*, 119(2):160–175, 1995.

[9] H.-D. Ebbinghaus and J. Flum. *Finite model theory.* Perspectives in Mathematical Logic. Springer, 1995.

[10] M. Fürer. The computational complexity of the unconstrained limited domino problem (with implications for logical decision problems). In E. Börger, G. Hasenjaeger, and D. Rödding, editors, *Logic and Machines: Decision Problems and Complexity, Proceedings of the Symposium "Rekursive Kombinatorik" held from May 23-28, 1983 at the Institut für Mathematische Logik und Grundlagenforschung der Universität Münster/Westfalen*, volume 171 of *Lecture Notes in Computer Science*, pages 312–319. Springer, 1984.

[11] M. Fürer. Weisfeiler-lehman refinement requires at least a linear number of iterations. In F. Orejas, P. G. Spirakis, and J. van Leeuwen, editors, *Automata, Languages and Programming, 28th International Colloquium, ICALP 2001, Crete, Greece, July 8-12, 2001, Proceedings*, volume 2076 of *Lecture Notes in Computer Science*, pages 322–333. Springer, 2001.

[12] E. Grädel, P. G. Kolaitis, and M. Y. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic*, 3(1):53–69, 1997.

[13] E. Grädel and M. Otto. On logics with two variables. *Theor. Comput. Sci.*, 224(1-2):73–113, 1999.

[14] M. Grohe. Finite variable logics in descriptive complexity theory. *Bulletin of Symbolic Logic*, 4(4):345–398, 1998.

[15] M. Grohe. Equivalence in finite-variable logics is complete for polynomial time. *Combinatorica*, 19(4):507–532, 1999.

[16] M. Grohe. Fixed-point definability and polynomial time on graphs with excluded minors. In *Proceedings of the 25th Annual IEEE Symposium on Logic in Computer Science, LICS 2010, 11-14 July 2010, Edinburgh, United Kingdom*, pages 179–188. IEEE Computer Society, 2010.

[17] M. Grohe. From polynomial time queries to graph structure theory. *Commun. ACM*, 54(6):104–112, 2011.

[18] M. Grohe and O. Verbitsky. Testing graph isomorphism in parallel by playing a game. In M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, editors, *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part I*, volume 4051 of *Lecture Notes in Computer Science*, pages 3–14. Springer, 2006.

[19] I. M. Hodkinson. Finite variable logics. *Bulletin of the EATCS*, 51:111–140, 1993.

[20] N. Immerman. Number of quantifiers is better than number of tape cells. *J. Comput. Syst. Sci.*, 22(3):384–406, 1981.

[21] N. Immerman. *Descriptive complexity.* Graduate texts in computer science. Springer, 1999.

[22] N. Immerman and E. Lander. Describing graphs: A first-order approach to graph canonization. In A. Selman, editor, *Complexity Theory Retrospective*, pages 59–81. Springer-Verlag, 1990.

[23] P. G. Kolaitis and J. Panttaja. On the complexity of existential pebble games. In M. Baaz and J. A. Makowsky, editors, *Computer Science Logic, 17th International Workshop, CSL 2003, 12th Annual Conference of the EACSL, and 8th Kurt Gödel Colloquium, KGC 2003, Vienna, Austria, August 25-30, 2003, Proceedings*, volume 2803 of *Lecture Notes in Computer Science*, pages 314–329. Springer, 2003.

[24] P. G. Kolaitis and M. Y. Vardi. On the expressive power of datalog: Tools and a case study. *J. Comput. Syst. Sci.*, 51(1):110–134, 1995.

[25] P. G. Kolaitis and M. Y. Vardi. A game-theoretic approach to constraint satisfaction. In H. A. Kautz and B. W. Porter, editors, *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on on Innovative Applications of Artificial Intelligence, July 30 - August 3, 2000, Austin, Texas, USA*, pages 175–181. AAAI Press / The MIT Press, 2000.

[26] L. Libkin. *Elements of Finite Model Theory.* Springer, 2004.

[27] O. Matz, N. Schweikardt, and W. Thomas. The monadic quantifier alternation hierarchy over grids and graphs. *Inf. Comput.*, 179(2):356–383, 2002.

[28] E. Pezzoli. Computational complexity of ehrenfeucht-fraïssé games on finite structures. In G. Gottlob, E. Grandjean, and K. Seyr, editors, *Computer Science Logic, 12th International Workshop, CSL '98, Annual Conference of the EACSL, Brno, Czech Republic, August 24-28, 1998, Proceedings*, volume 1584 of *Lecture Notes in Computer Science*, pages 159–170. Springer, 1998.

[29] O. Pikhurko, J. Spencer, and O. Verbitsky. Succinct definitions in the first order theory of graphs. *Ann. Pure Appl. Logic*, 139(1-3):74–109, 2006.

[30] O. Pikhurko and O. Verbitsky. Logical complexity of graphs: a survey. In M. Grohe and J. Makowsky, editors, *Model theoretic methods in finite combinatorics*, volume 558 of *Contemporary Mathematics*, pages 129–179. American Mathematical Society, Providence, RI, 2011.

[31] M. Y. Vardi. On the complexity of bounded-variable queries. In M. Yannakakis, editor, *Proceedings of the Fourteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, May 22-25, 1995, San Jose, California, USA*, pages 266–276. ACM Press, 1995.

[32] P. Weis and N. Immerman. Structure theorem and strict alternation hierarchy for FO^2 on words. *Logical Methods in Computer Science*, 5(3), 2009.